

# **Virtio ring and virtio-net**

Michael S Tsirkin

Red Hat

news

case study

challenges

# Virtio with kvm

**virtio** Guest

**KVM**

**eventfd**

**Vhost**

**qemu**

**io**

**virtio**

**Guest**

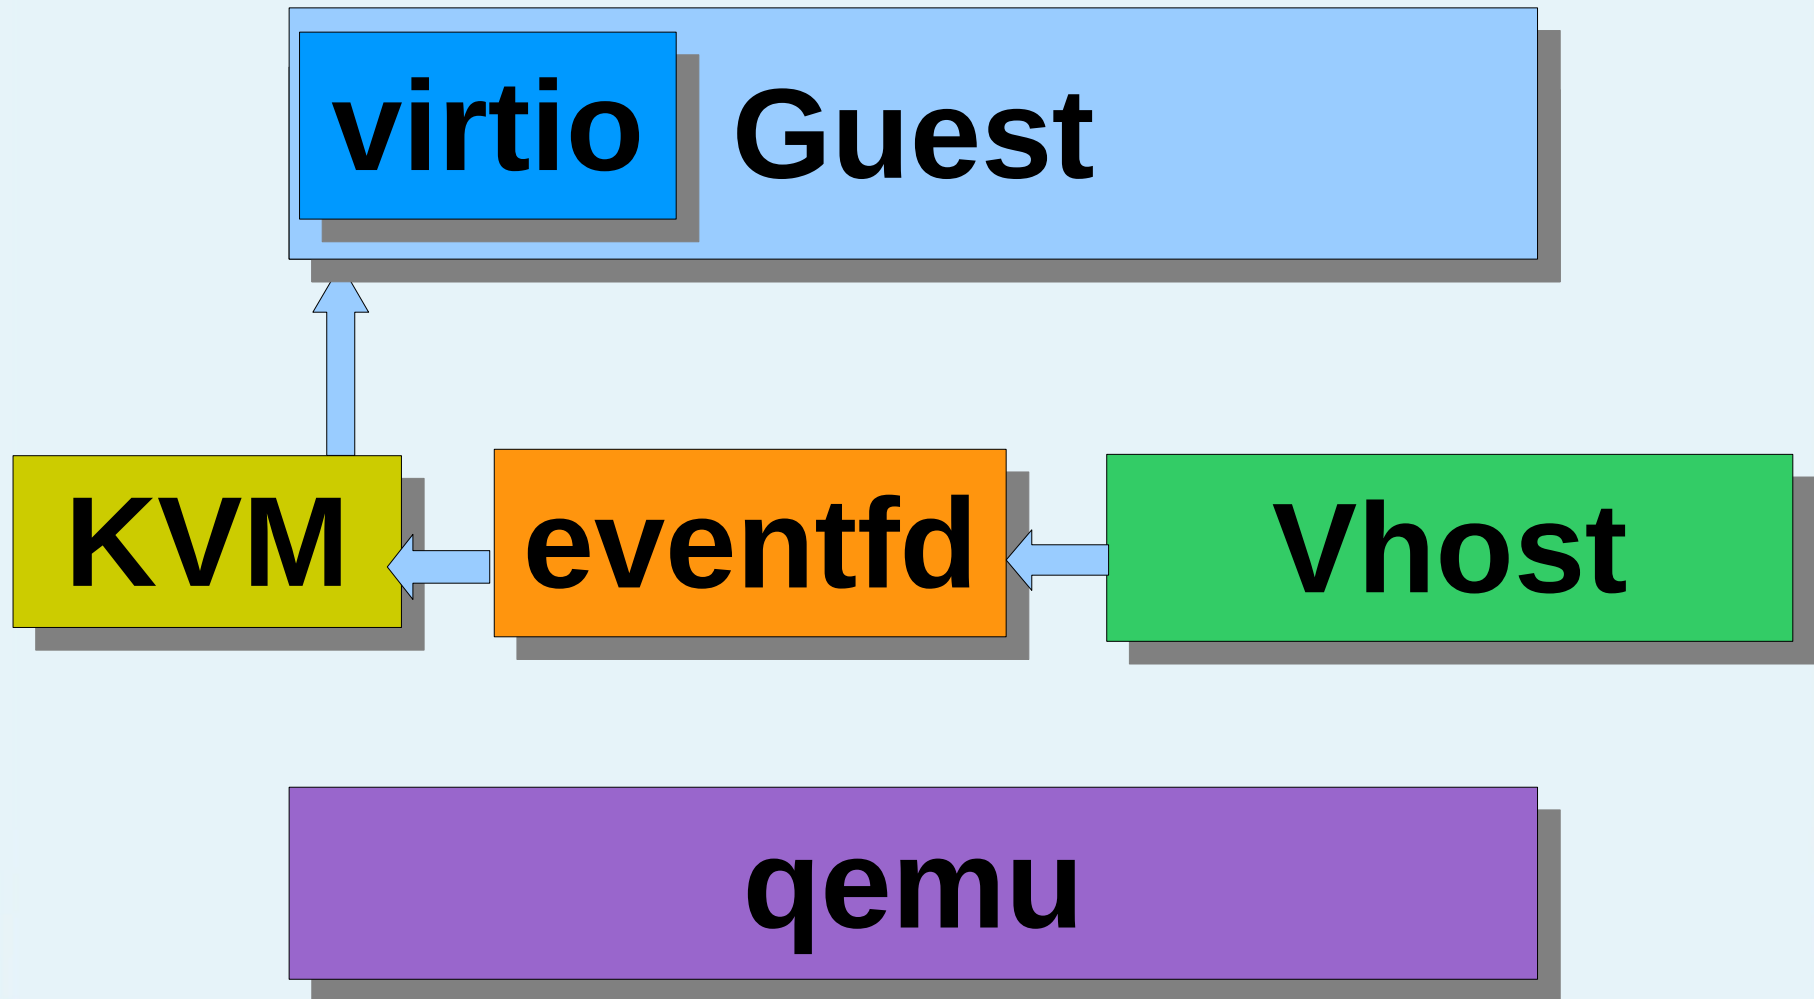
**KVM**

**eventfd**

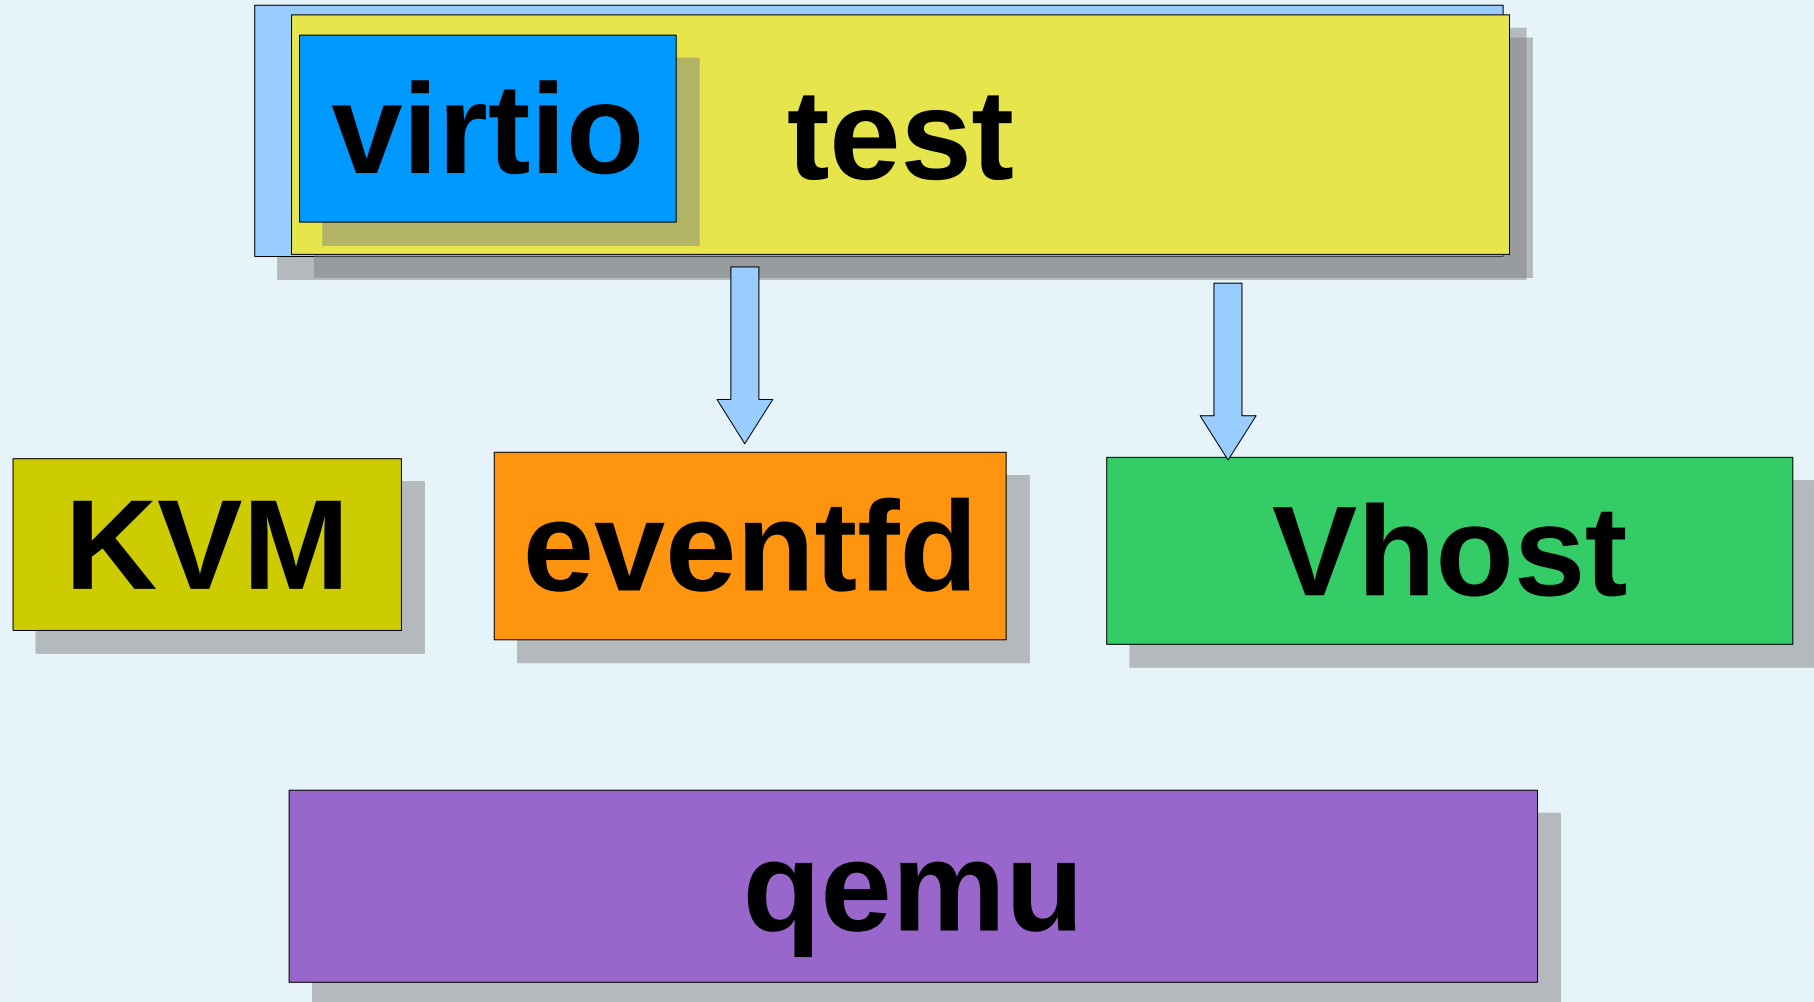
**Vhost**

**qemu**

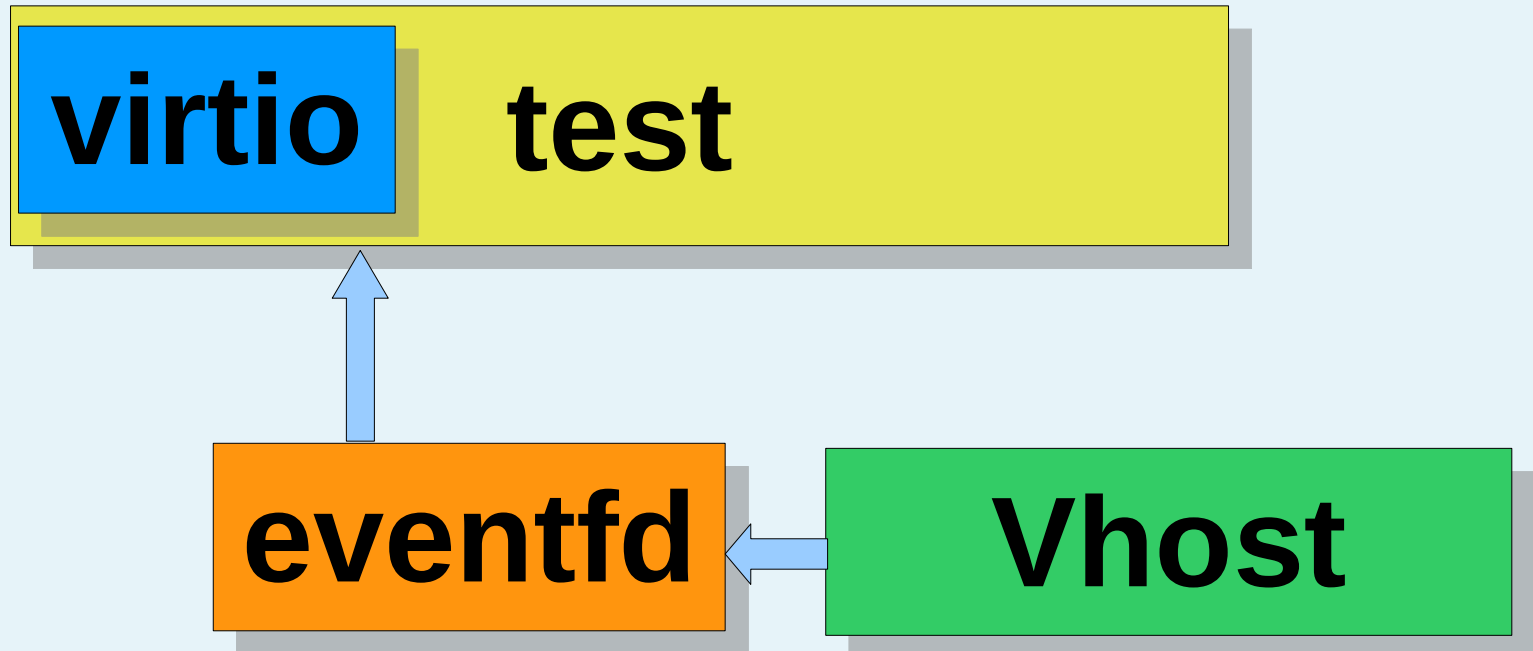
# interrupt



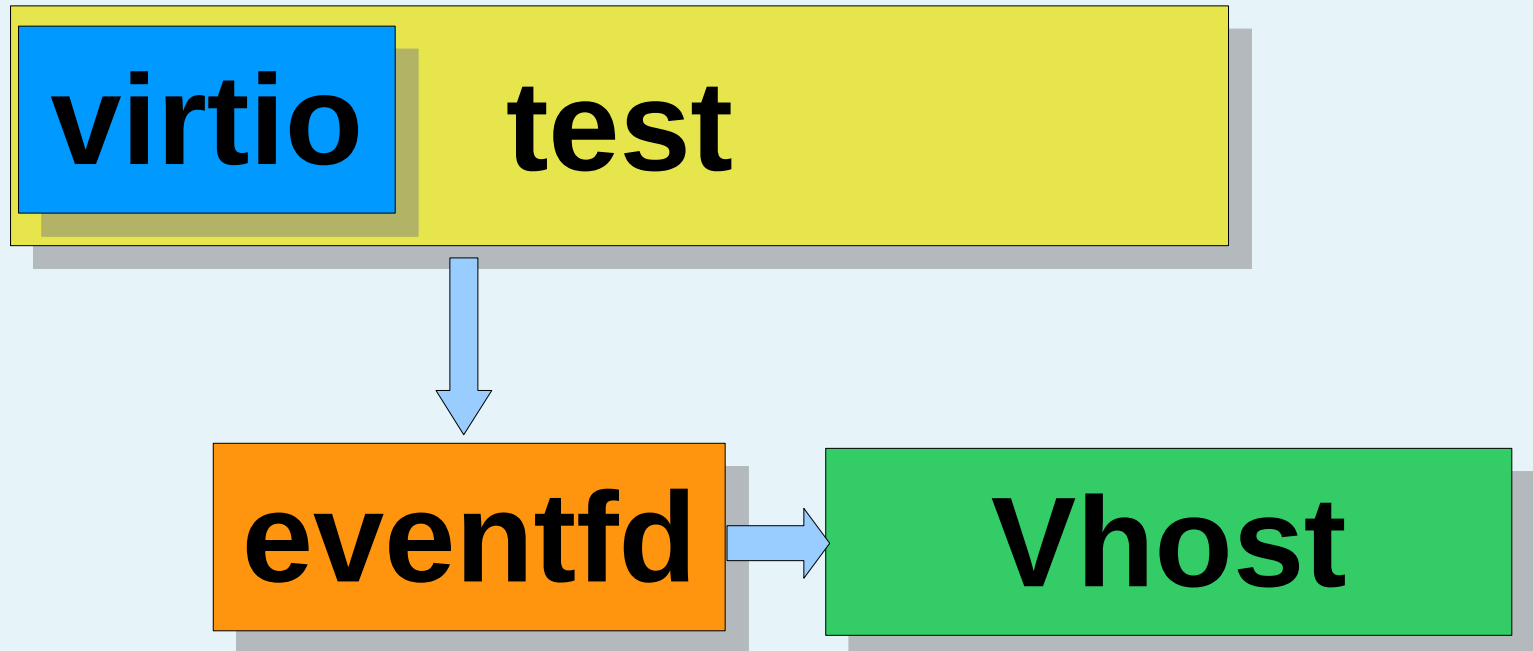
# virtio-test



# poll



# **write**

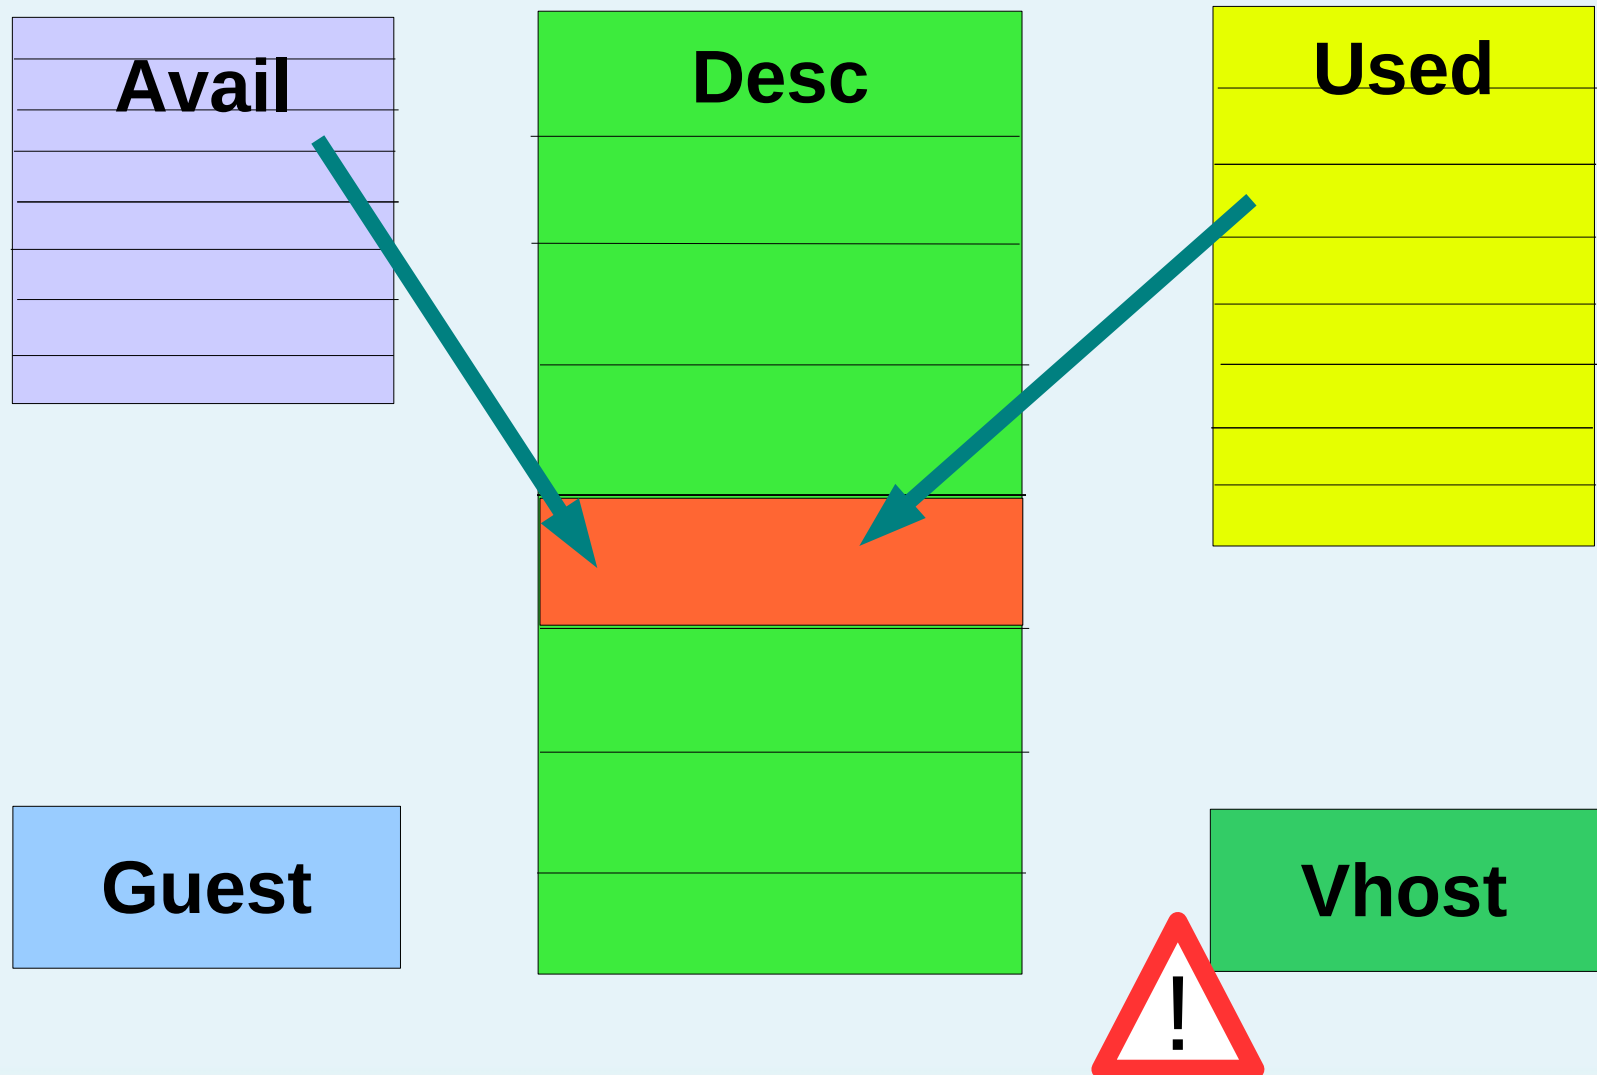


# linux/tools/virtio


```
# vi drivers/virtio/virtio.c
....
# cd tools/virtio/ ; make -q
# insmod vhost_test/vhost_test.ko
# time ./virtio_test
real      0m0.649s
user      0m0.295s
sys       0m0.335s
```




# Virtio ring



# window



# Event index



# Extending virtio



**test**

**virtio**

**Guest**

**KVM**

**eventfd**

**Vhost**

**qemu**

# linux/tools/virtio

```
# time ./virtio_test --no-event-idx
```

```
real      0m0.649s
```

```
user      0m0.295s
```

```
sys       0m0.335s
```

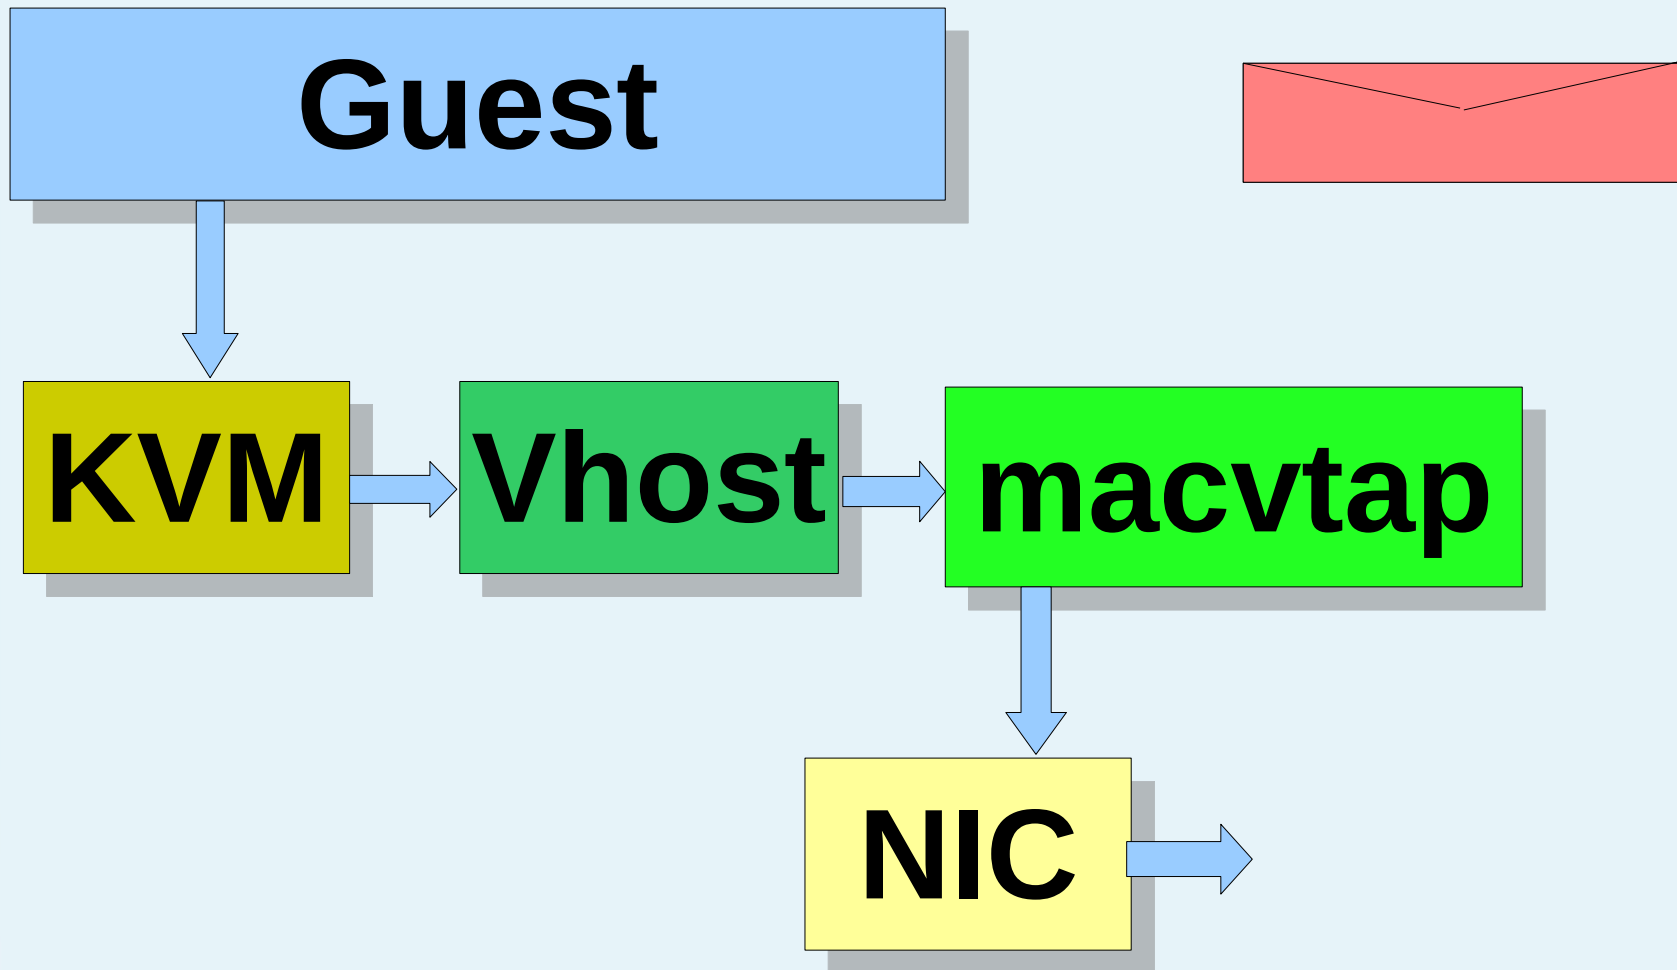
```
# time ./virtio_test --event-idx
```

```
real      0m0.169s
```

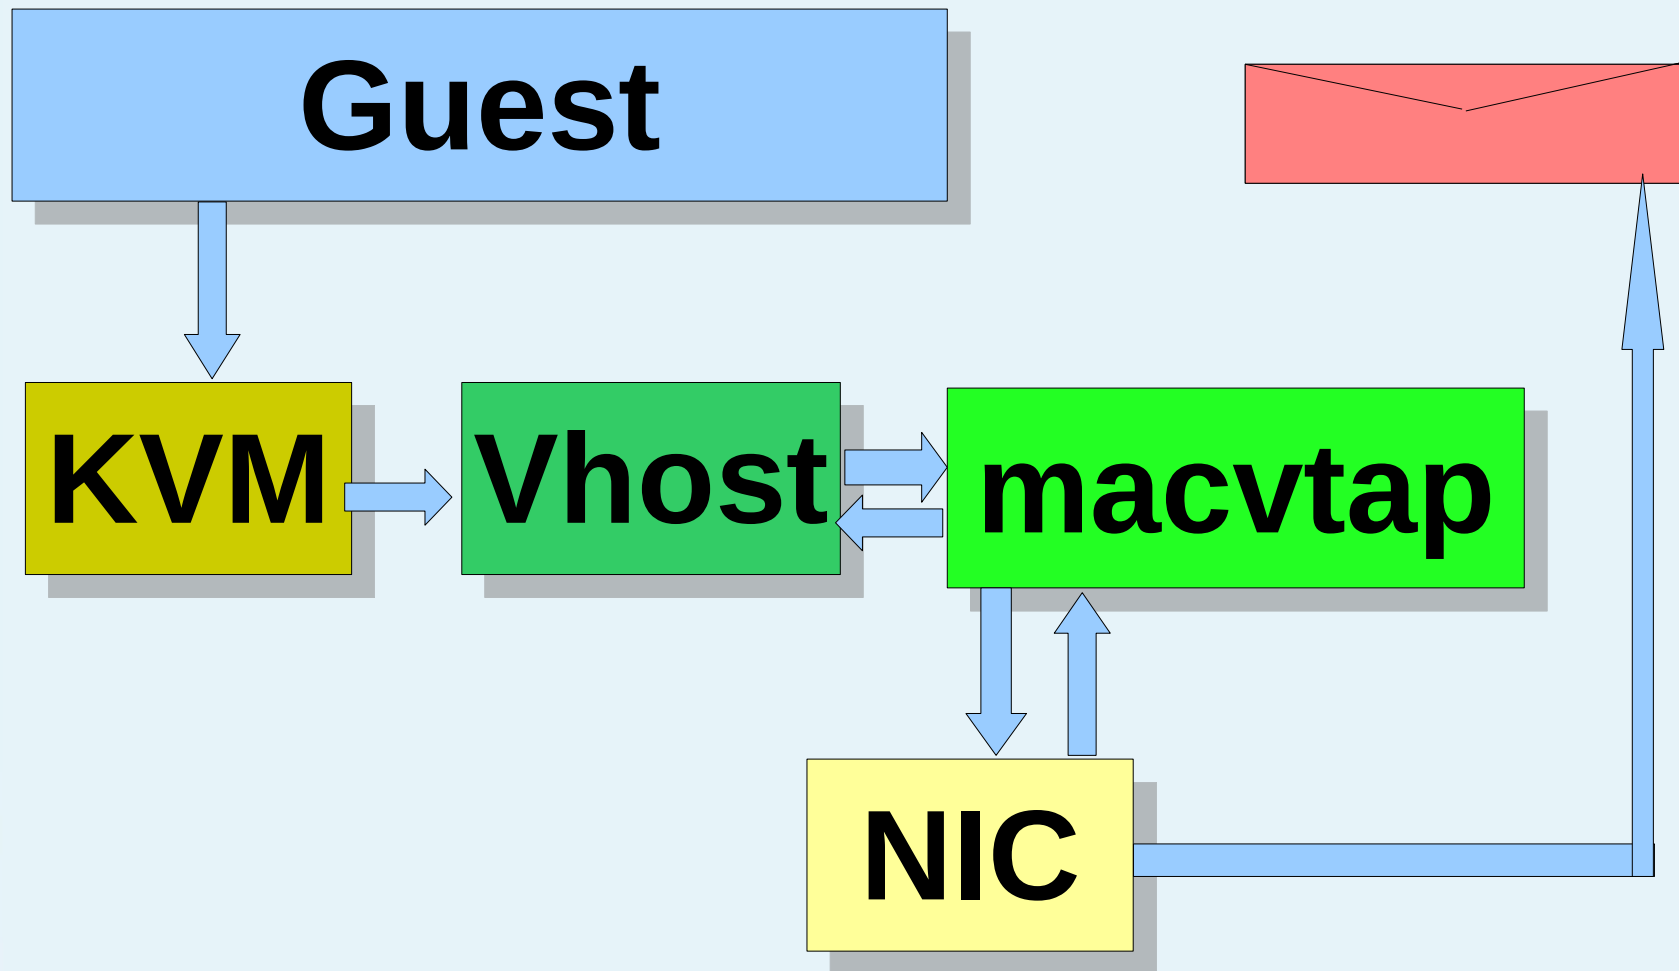
```
user      0m0.140s
```

```
sys       0m0.019s
```

# virtio-net TX




# Zcopy TX



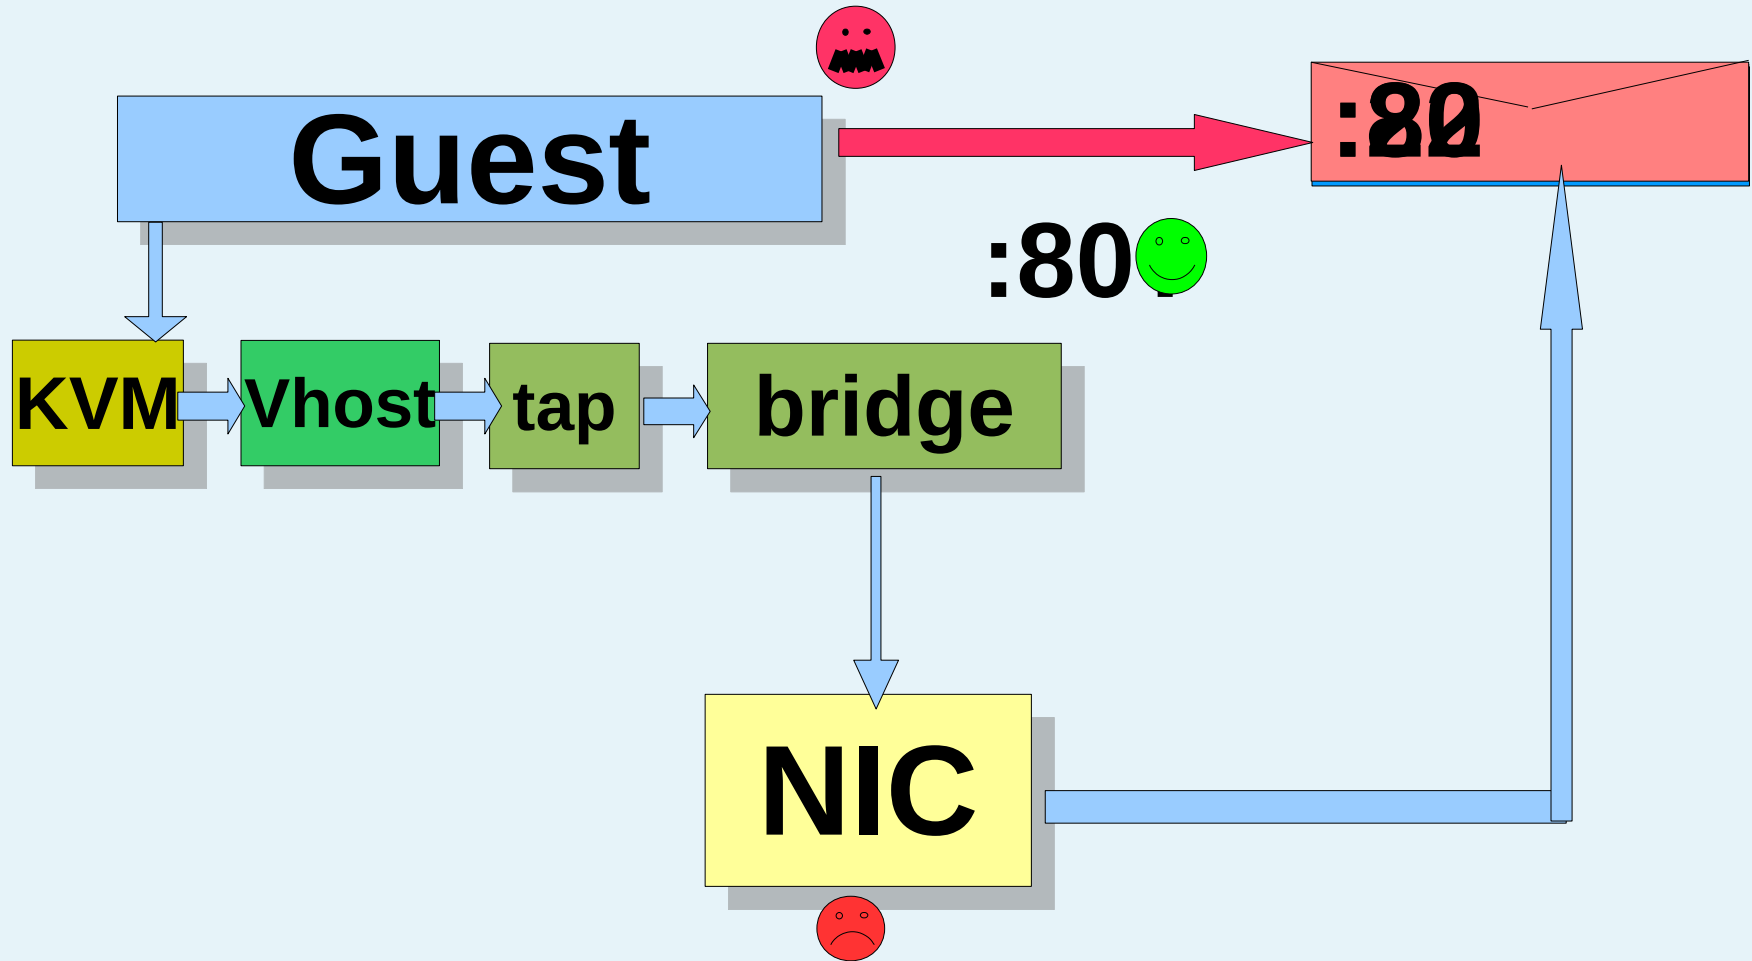
# Guest -> External

10 Gb/s

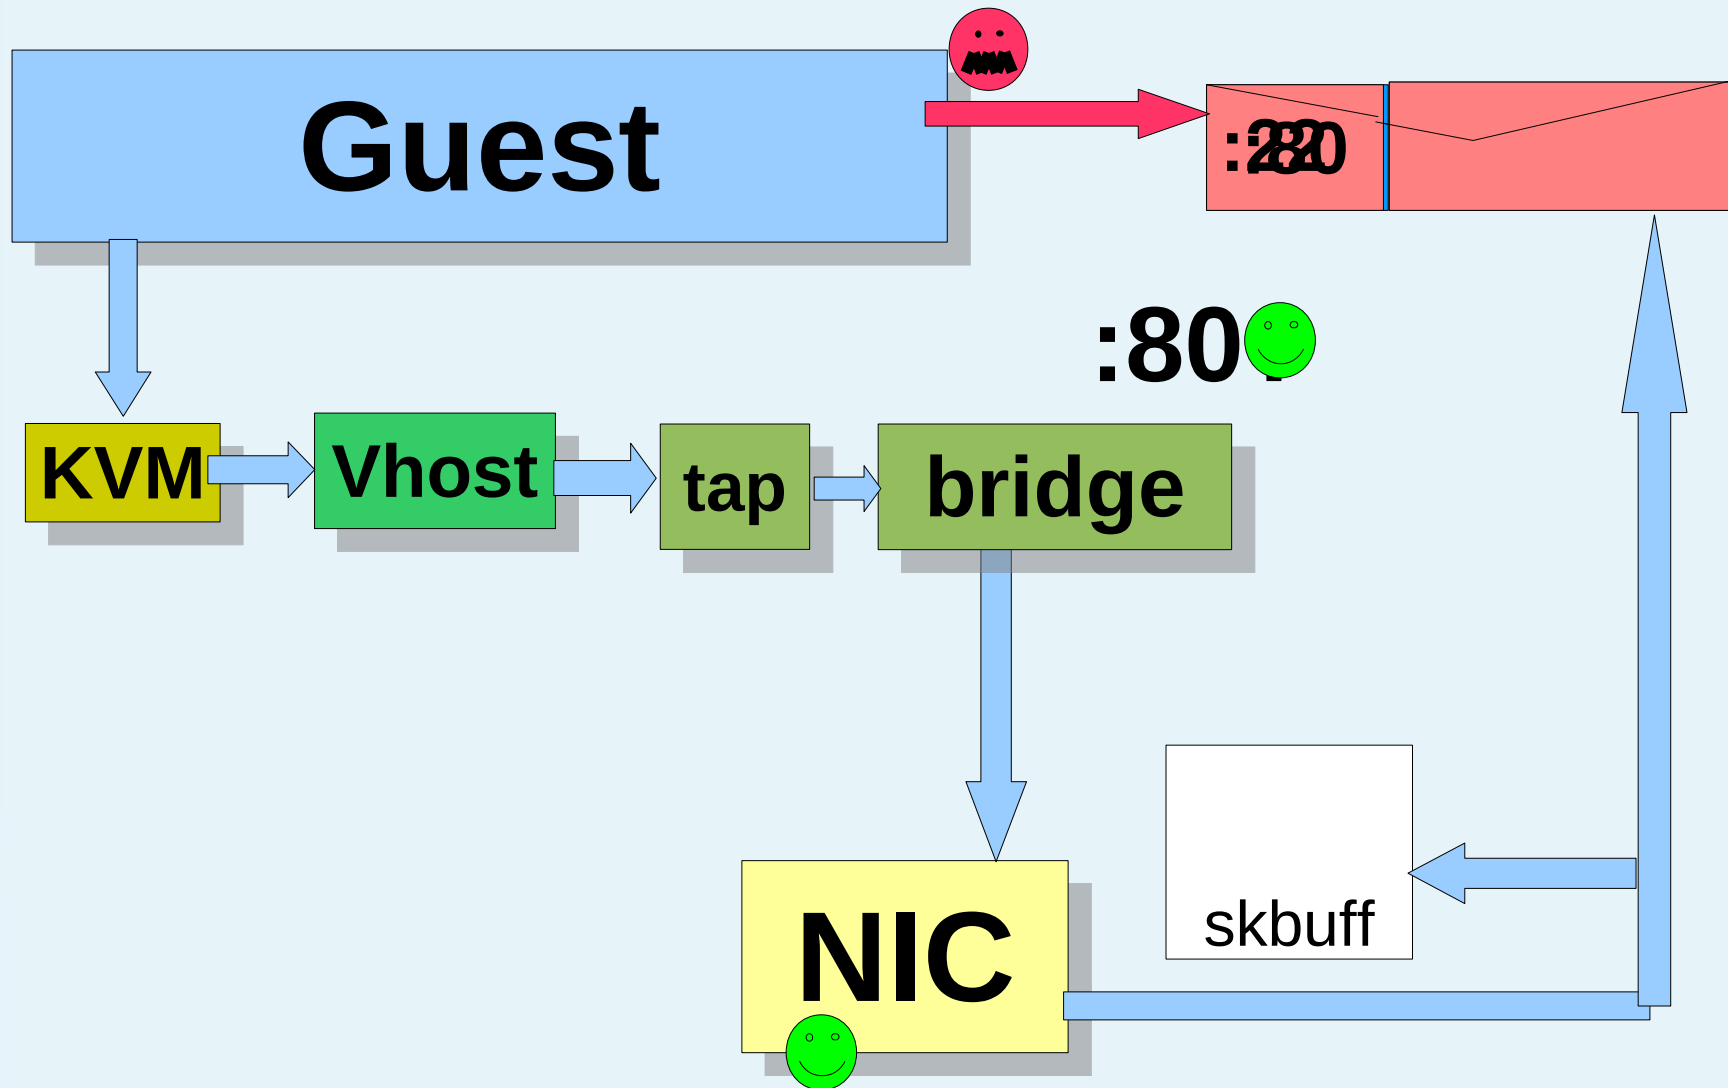




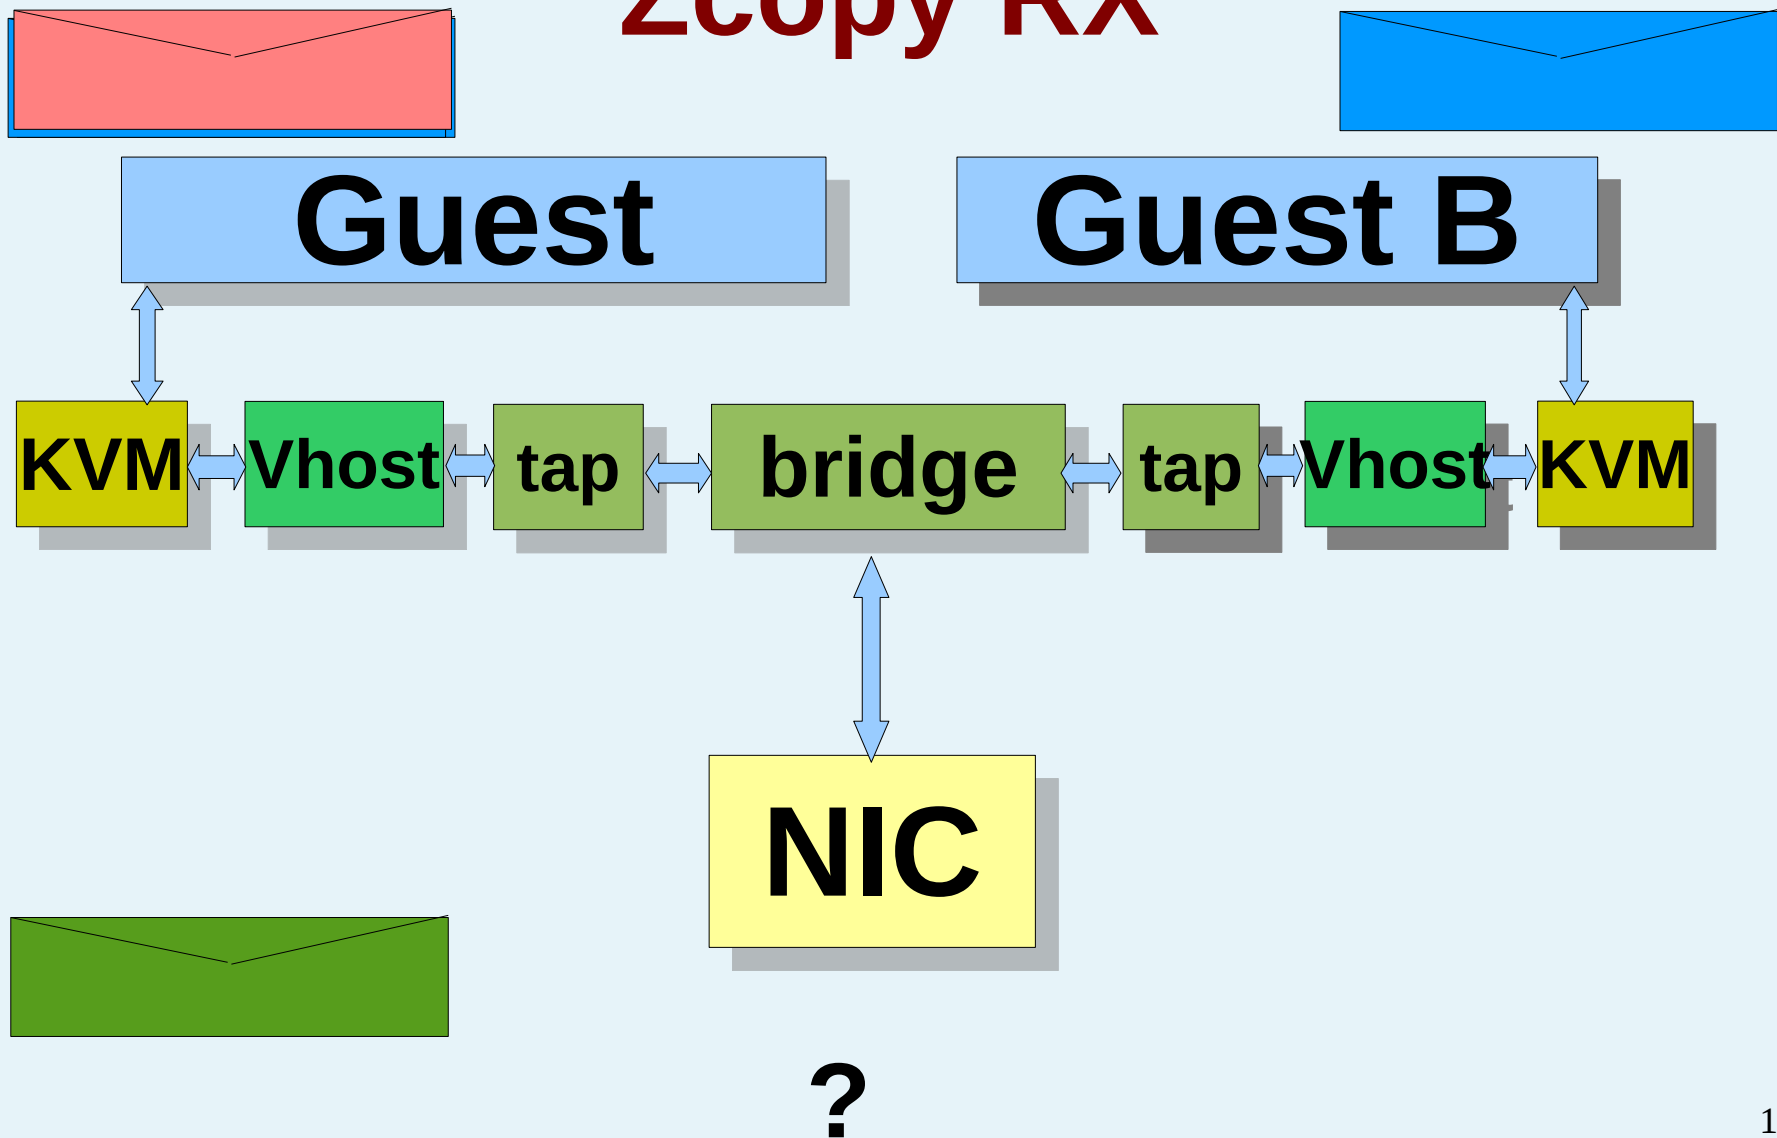
# Filtering



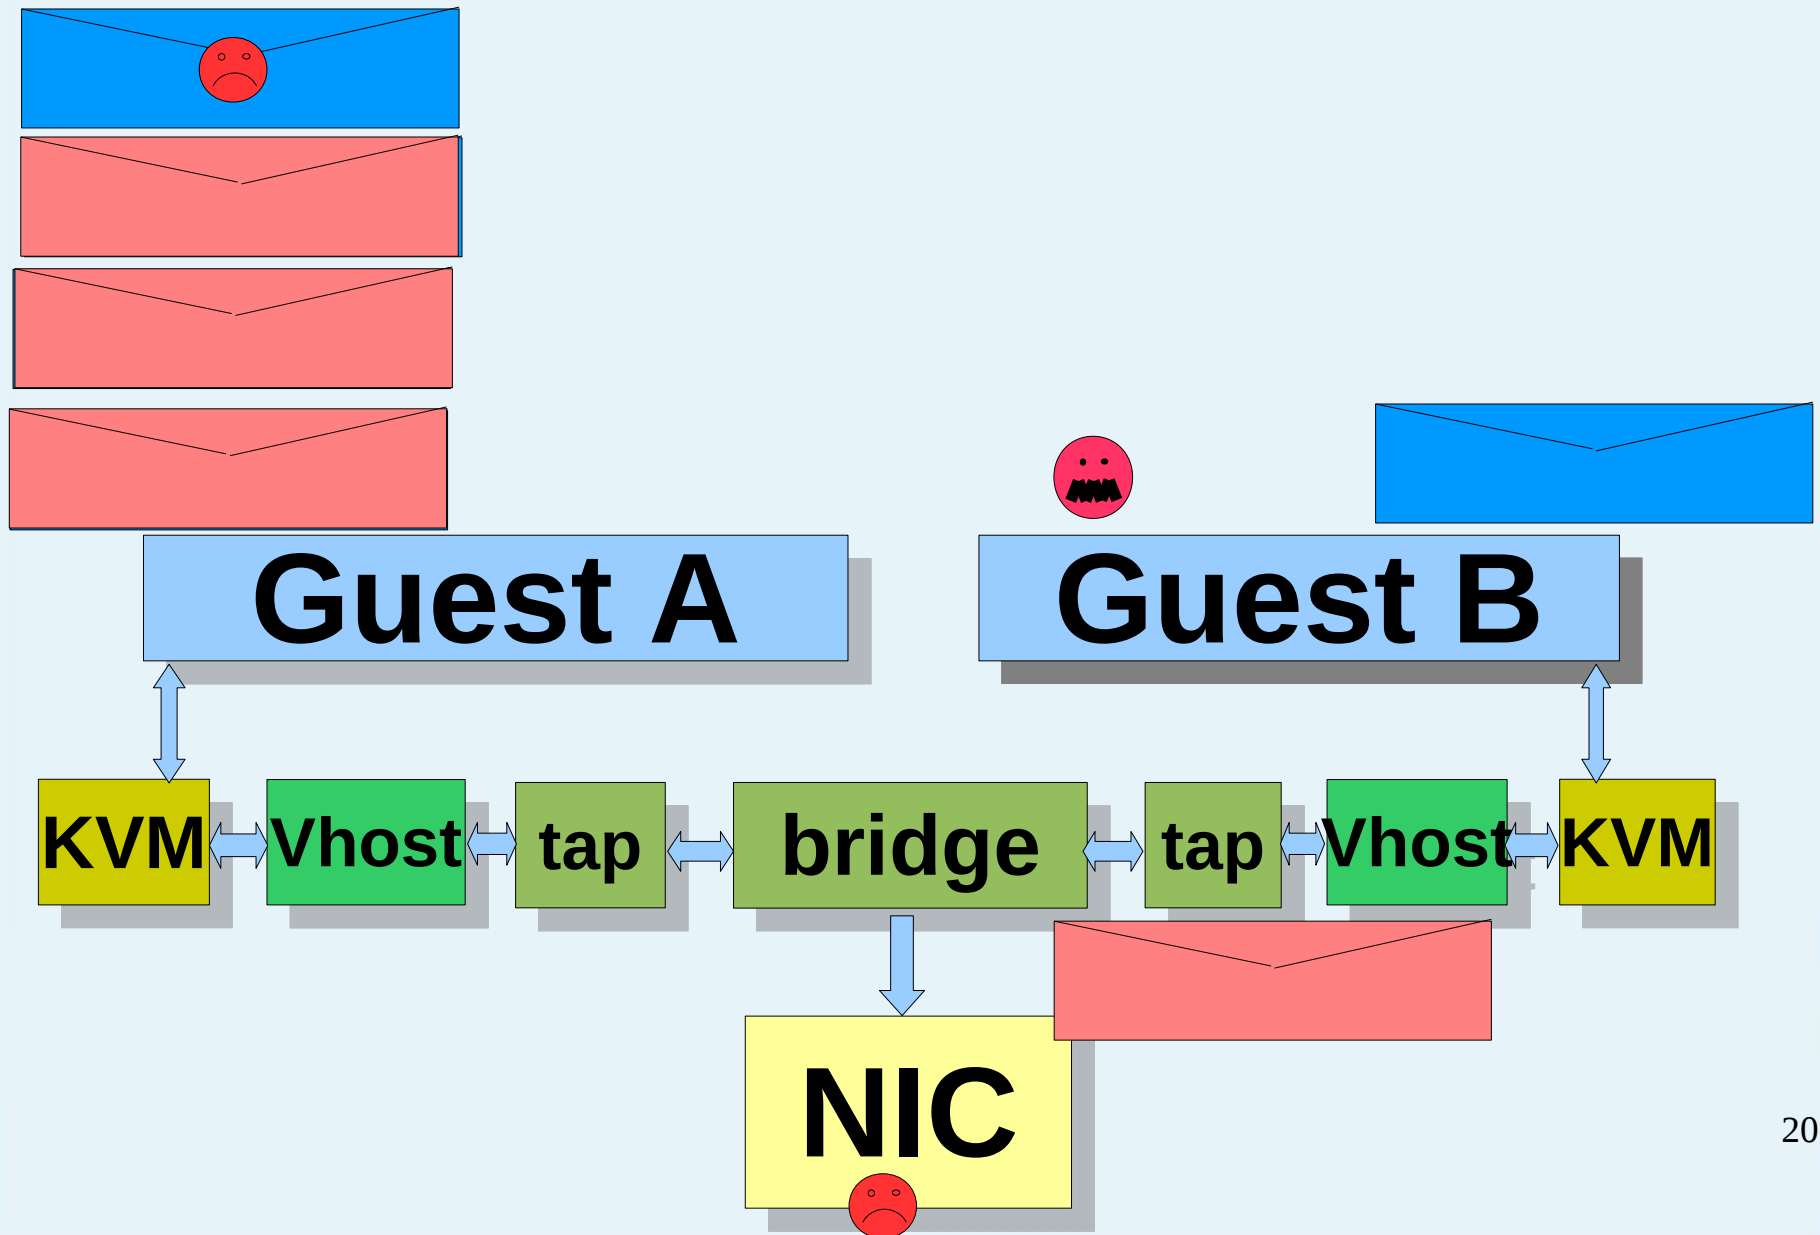
# Partial copy



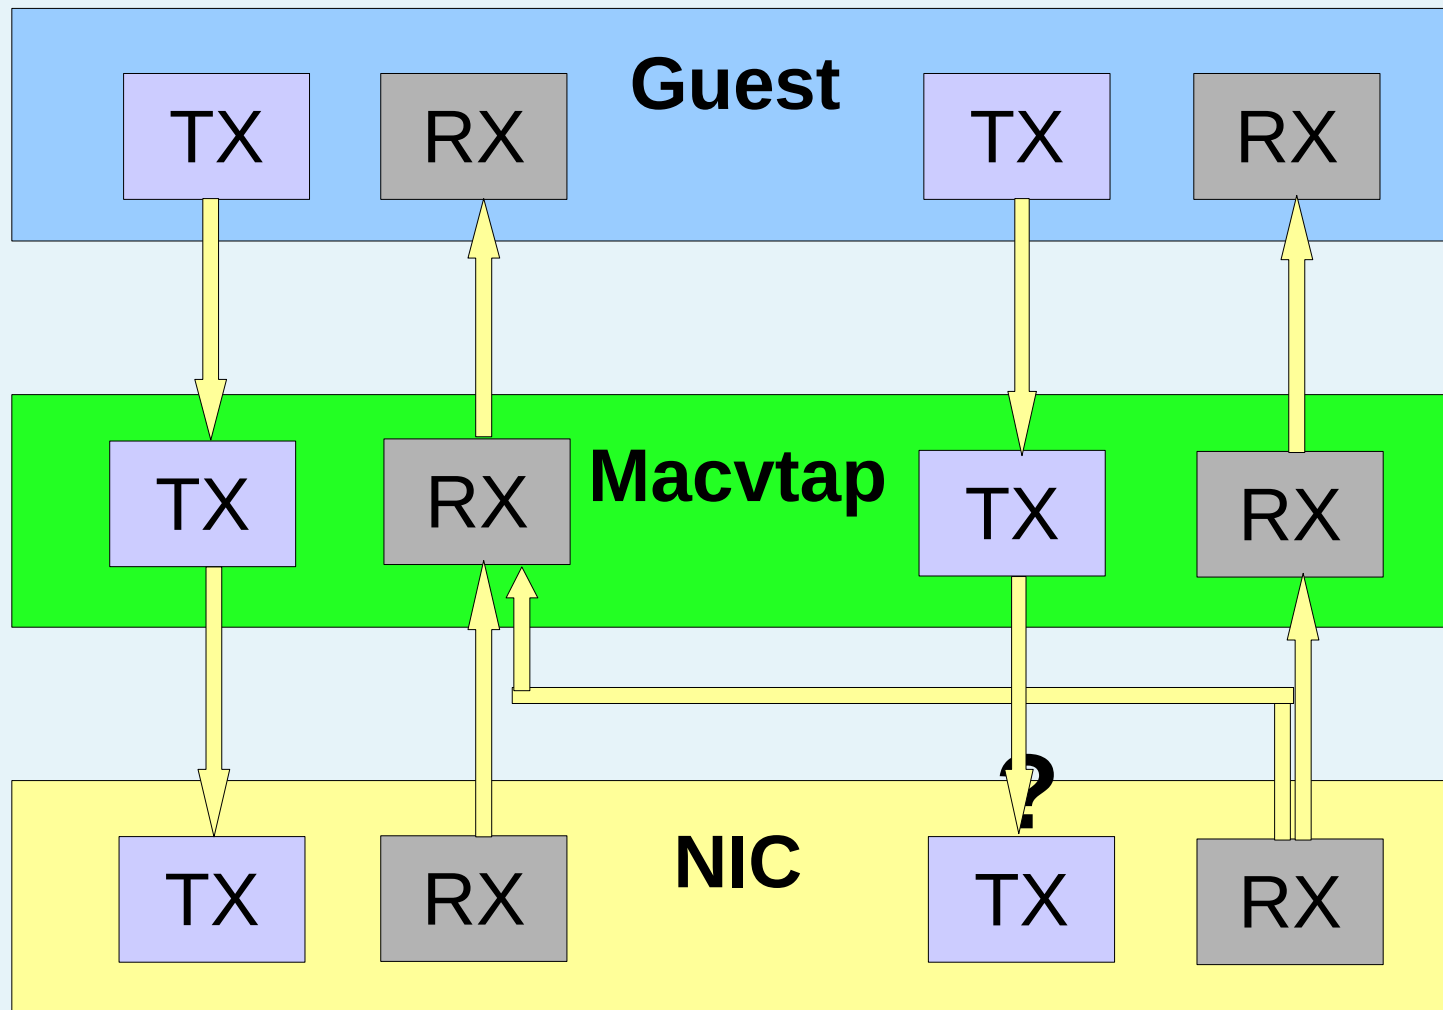
# Zcopy RX



# Guest 2 Guest





# multiqueue



# External -> Guest

10 Gb/s



# Status/TODO

- **Virtio ring** +
- **Zero copy TX** +
- **Multiqueue** +
- **Zero copy RX** ?
- **Bridge zero copy** ?
- **Level interrupts** ?

# Questions?

-



# Virtio ring

